arXiv:1607.04527v1 [cs.DS] 15 Jul 2016

Maximizing a Monotone Submodular Function with a Bounded Curvature under a Knapsack Constraint

Yuichi Yoshida*
National Institute of Informatics *and* Preferred Infrastructure, Inc.
yyoshida@nii.ac.jp

July 18, 2016

Abstract

We consider the problem of maximizing a monotone submodular function under a knapsack constraint. We show that, for any fixed $\epsilon > 0$, there exists a polynomial-time algorithm with an approximation ratio $1 - c/e - \epsilon$, where $c \in [0, 1]$ is the (total) curvature of the input function. This approximation ratio is tight up to ϵ for any $c \in [0, 1]$. To the best of our knowledge, this is the first result for a knapsack constraint that incorporates the curvature to obtain an approximation ratio better than $1 - 1/e$, which is tight for general submodular functions.

As an application of our result, we present a polynomial-time algorithm for the budget allocation problem with an improved approximation ratio.

*Supported by JSPS Grant-in-Aid for Young Scientists (B) (26730009), MEXT Grant-in-Aid for Scientific Research on Innovative Areas (24106003), and JST, ERATO, Kawarabayashi Large Graph Project.

1 Introduction

In this paper, we consider the problem of maximizing a monotone submodular function under a knapsack constraint. Specifically, given a monotone submodular function $f : 2^E \to \mathbb{R}_+$ and a weight function $w : E \to [0, 1]$, we aim to solve the following optimization problem:

$$\text{maximize } f(S) \quad \text{subject to } w(S) \leq 1 \text{ and } S \subseteq E,$$

where $w(S) = \sum_{e \in S} w(e)$. This problem has wide applications in machine learning tasks such as sensor placement [8], document summarization [11, 12], maximum entropy sampling [10], and budget allocation [14]. Although this problem is NP-hard in general, it is known that we can achieve $(1 - 1/e)$-approximation in polynomial time [15], and this approximation ratio is indeed tight [5].

Although it is useful to know that we can always obtain $(1 - 1/e)$-approximation in polynomial time, it is observed that a simple greedy method outputs even better solutions in real applications (see, e.g., [8]), and it is more desirable if we can guarantee a better approximation ratio by making assumptions on the input function. One such assumption is the notion of curvature, introduced by Conforti and Cornuéjols [4]. For a monotone submodular function $f : 2^E \to \mathbb{R}_+$, the *(total) curvature* of f is defined as

$$c_f = 1 - \min_{e \in E} \frac{f_{E-e}(e)}{f(e)}.$$

Intuitively speaking, the curvature measures how close f is to a linear function. To see this, note that $c_f \in [0, 1]$ and $c_f = 0$ if and only if f is a linear function.

It was shown in [4] that, for maximizing a monotone submodular function under a cardinality constraint, the greedy algorithm achieves an approximation ratio $(1 - e^{-c_f})/c_f$, and the result was extended to a matroid constraint [17]. Recently, Sviridenko *et al.* [16] obtained a polynomial-time algorithm for a matroid constraint with an approximation ratio $1 - c_f/e$, and showed that this approximation ratio is indeed tight for every $c_f \in [0, 1]$ even under a cardinality constraint (note that $1 - c_f/e$ is strictly larger than $(1 - e^{-c_f})/c_f$ except when $c_f = 0$ or $c_f = 1$).

In this paper, we extend these results to a knapsack constraint and present a polynomial-time algorithm under a knapsack constraint with an approximation ratio $1 - c_f/e$. More specifically, we show the following:

Theorem 1.1. *There is an algorithm that, given a monotone submodular function $f : 2^E \to \mathbb{R}_+$, a weight function $w : E \to [0, 1]$, and $\epsilon \in (0, 1)$, outputs a (random) set $S \subseteq E$ with $w(S) \leq 1$ (with probability one) satisfying*

$$\mathbf{E}[f(S)] \geq \left(1 - \frac{c_f}{e} - \epsilon\right) f(O).$$

Here, $O \subseteq E$ is an optimal solution to the problem, i.e., O is a set with $w(O) \leq 1$ that maximizes f. The running time is $O\left(n^5 + n^4 \mathrm{polylog}(n) \cdot \left(\frac{1}{\epsilon}\right)^{\mathrm{poly}(1/\epsilon)}\right)$, where $n = |E|$.

Note that the approximation ratio $1 - c_f/e$ is indeed tight for every $c_f \in [0, 1]$ because the lower bound given by [16] holds even for a cardinality constraint. To the best of our knowledge, this is the first result for a knapsack constraint that incorporates the curvature to obtain an approximation ratio better than $1 - 1/e$, which is tight for general submodular functions.

We can apply our algorithm to all the above-mentioned applications to obtain a better solution when the input function has a small curvature. As a representative example, we consider the

budget allocation problem [1], which models a marketing process that allocates a given budget among media channels, such as TV, newspapers, and the Web, in order to maximize the impact on customers. We model the process using a bipartite graph on a vertex set $A \cup B$, where A and B correspond to the media channels and the customers, respectively, and an edge $(a, b) \in A \times B$ represents the potential influence of media channel a on customer b. In a simplified setting where we can use each channel at most once, each media channel $a \in A$ can activate a customer with a predetermined probability $p_a \in [0, 1]$. Then, we have to find a set $S \subseteq A$ that maximizes the expected number of activated customers subject to $w(S) := \sum_{a \in S} w(a) \leq 1$, where $w(a)$ is the cost of using media channel a. We can formulate this problem as the maximization of a monotone submodular function $f : 2^A \to \mathbb{R}_+$ under the knapsack constraint $w(S) \leq 1$. We can show that $c_f \leq 1 - \min_{b \in B} p^{|\Gamma(b)|-1}$, where $\Gamma(b)$ denotes the set of neighbors of b in the bipartite graph. By Theorem 1.1, this immediately gives the approximation ratio $1 - c_f/e - \epsilon$ for this problem. The actual model is more general and discussed in detail in Section 5.

1.1 Proof technique

Now, we present the outline of our proof. Let $f : 2^E \to \mathbb{R}_+$ be the input function and $O \subseteq E$ be the optimal solution, i.e., O is the set that maximizes f among the sets with weight at most one. We assume that $c_f = 1 - \Omega(\epsilon)$; otherwise, we can use a standard algorithm [15] to achieve the desired approximation ratio. Using the argument in [16], we can decompose the input function f into a monotone submodular function $g : 2^E \to \mathbb{R}_+$ and a linear function $\ell : 2^E \to \mathbb{R}_+$ such that, if we can compute a set $S \subseteq E$ with $w(S) \leq 1$ and $f(S) = g(S) + \ell(S) \geq (1 - 1/e)g(O) + \ell(O)$, then S is a $(1 - c_f/e)$-approximate solution. Moreover, by slightly changing the argument in [16], we can also assume that $c_g = 1 - \Omega(\epsilon(1 - c_f)) = 1 - \Omega(\epsilon^2)$.

In order to find the desired set $S \subseteq E$, we use a variant of the continuous greedy algorithm [3] that simultaneously optimizes g and ℓ. In this algorithm, we consider continuous versions of g, ℓ, and w, denoted by $G : [0, 1]^E \to \mathbb{R}_+$, $L : [0, 1]^E \to \mathbb{R}_+$, and $W : [0, 1]^E \to \mathbb{R}_+$, respectively. We note that the function G is called the multilinear extension of g, and that L and W are linear functions. We start with the zero vector $\boldsymbol{x} \in [0, 1]^E$ and then iteratively update it. The algorithm consists of $1/\epsilon$ iterations, and roughly speaking, in each iteration, we find a vector $\boldsymbol{v} \in [0, 1]^E$ with the following properties: (i) $G(\boldsymbol{x} + \epsilon\boldsymbol{v}) - G(\boldsymbol{x}) \geq \epsilon(G(\boldsymbol{x} \vee \mathbf{1}_O) - G(\boldsymbol{x}))$, (ii) $L(\epsilon\boldsymbol{v}) = \epsilon L(\boldsymbol{v}) \geq \epsilon\ell(O)$, and (iii) $W(\epsilon\boldsymbol{v}) = \epsilon W(\boldsymbol{v}) \leq w(O)$. Then, we update $\boldsymbol{x}$ by adding $\epsilon\boldsymbol{v}$. Here, $\mathbf{1}_O$ is the characteristic vector of the set O and $\vee$ is the coordinate-wise maximum. Intuitively speaking, these conditions mean that moving along the direction $\boldsymbol{v}$ from $\boldsymbol{x}$ is no worse than moving towards $\boldsymbol{x} \vee \mathbf{1}_O$. We can find such a vector $\boldsymbol{v}$ by linear programming. Then, after $1/\epsilon$ iterations, we get a vector $\boldsymbol{x} \in [0, 1]^E$ such that $G(\boldsymbol{x}) \geq (1 - 1/e)G(\mathbf{1}_O) = (1 - 1/e)g(O)$, $L(\boldsymbol{x}) \geq \ell(O)$, and $W(\boldsymbol{x}) \leq w(O)$. Finally, we obtain a set $S \subseteq E$ by rounding the vector $\boldsymbol{x}$, where each element $e \in E$ is added with probability $\boldsymbol{x}(e)$.

Unfortunately, this strategy does not work as is. Here, a crucial issue is that we cannot show the concentration of the weight in the rounding step. To address this issue, by borrowing an idea from [2], we split the elements into large and small ones, where an element is said to be *small* if $g(e) \leq \epsilon^6 g(O)$ and $\ell(e) \leq \epsilon^6 \ell(O)$, and is said to be *large* otherwise (in our analysis, it is more convenient to define large and small elements in terms of g and ℓ instead of w). Then, since the curvature of g is bounded away from one, we can bound the number of large elements in O by a

function of ϵ.[1] Let $O_{\mathrm{L}}, O_{\mathrm{S}} \subseteq O$ be the set of large and small elements in O, respectively. Further, we let $O_{\mathrm{L}} = \{o_1, \dots, o_m\}$. Then, in each iteration, we do the following: For each $i \in \{1, \dots, m\}$, we find an element e_i such that (i) $G(\boldsymbol{x} \vee \mathbf{1}_{e_i}) - G(\boldsymbol{x}) \geq G(\boldsymbol{x} \vee \mathbf{1}_{o_i}) - G(\boldsymbol{x})$, (ii) $\ell(e_i) \geq \ell(o_i)$, and (iii) $w(e_i) \leq w(o_i)$. Then, we update $\boldsymbol{x}$ by adding $\epsilon \mathbf{1}_{e_i}$. Here, $\mathbf{1}_e$ is a characteristic vector of the element $e \in E$. Intuitively speaking, adding e_i to the current solution is no worse than adding o_i. For small items, we find a vector $\boldsymbol{v}$ as before by considering the characteristic vector $\mathbf{1}_{O_{\mathrm{S}}}$; then, we update $\boldsymbol{x}$ by adding $\epsilon \boldsymbol{v}$. In the rounding step, we handle large and small elements separately. Note that, for each $i \in \{1, \dots, m\}$, we have computed $1/\epsilon$ elements (through $1/\epsilon$ iterations). Then, we chose one of them uniformly at random and add it to the output set. An advantage of this rounding procedure is that we can guarantee that the chosen element for $i \in \{1, \dots, m\}$ has weight at most $w(o_i)$. For small elements, we apply the previous rounding procedure with a minor tweak to guarantee that the output set has weight at most one.

In order to realize this idea, we need to address several additional issues. First, as we do not know the set O, we do not know values related to O, such as $g(O)$, $\ell(O)$, $G(\boldsymbol{x} \vee \mathbf{1}_O)$, and $G(\boldsymbol{x} \vee \mathbf{1}_{o_i})$. Hence, we cannot determine whether an element is small or large, and we cannot find the desired vector or element in each iteration. We address this issue by guessing these values. For example, we can show a lower bound and an upper bound on $g(O)$ that are $O(n)$ times apart. This means that we can find a $(1-\epsilon)$-approximation to $g(O)$ in the geometric sequence of length $O(\log_{1+\epsilon} n) = O((\log n)/\epsilon)$ between the lower and upper bounds. If we naively guess all the values, as we have $1/\epsilon$ iterations, the resulting time complexity will be $\mathrm{poly}(n) \cdot \big((\log n)/\epsilon\big)^{\mathrm{poly}(1/\epsilon)}$. However, since the function g has curvature $1 - \Omega(\epsilon^2)$, we can reduce the number of candidate values and thus improve the time complexity to $\mathrm{poly}(n) \cdot (1/\epsilon)^{\mathrm{poly}(1/\epsilon)}$.

1.2 Related work

As mentioned earlier, it has been shown that the greedy method achieves $(1 - e^{-c_f})/c_f$ approximation for a cardinality constraint [4]. The result was extended to a matroid constraint by Vondrák [17]. He showed that the result actually holds if we replace c_f with the *curvature to the optimum* c_f^*, and the approximation ratio $(1 - e^{-c_f^*})/c_f^*$ is tight. Sviridenko *et al.* [16] improved the approximation ratio to $1 - c_f/e - \epsilon$ for a matroid constraint (and hence, a cardinality constraint), which is unattainable with c_f^*, and showed that the approximation ratio $1 - c_f/e$ is tight even for a cardinality constraint.

Curvature has been used to explain the empirical performance of the greedy method. Sharma *et al.* [13] considered maximum entropy sampling on Gaussian radial basis functions (RBF) kernels, which can be modeled as the maximization of a monotone submodular function, and showed that the curvature of this problem is close to zero.

The maximization of a submodular function under a knapsack constraint has been studied extensively. Sviridenko obtained a $(1-1/e)$-approximation algorithm with time complexity $O(n^5)$ [15]. We can also obtain $(1 - 1/e - \epsilon)$-approximation with a constant number of knapsack constraints; however, the time complexity blows up to $n^{\mathrm{poly}(1/\epsilon)}$ [9]. It has been claimed in [2] that, for any fixed $\epsilon > 0$, there is a $(1-1/e-\epsilon)$-approximation algorithm with time complexity $\tilde{O}(n^2)$. However, as mentioned in the footnote, their argument has a drawback. Several approximation guarantees have been achieved in [7] using various parameters of the input function. However, none of them

[1] Although it is claimed in [2] that the number of large elements is bounded for any submodular function, it is not true in general.

has an approximation ratio better than $1 - 1/e$ based solely on the assumption that the curvature is bounded.

1.3 Organization

The remainder of this paper is organized as follows. Section 2 introduces the definitions used throughout the paper and reviews the basic properties of submodular functions. Section 3 explains the reduction to a joint approximation of a monotone submodular function and a monotone linear function. Section 4 presents a joint approximation algorithm. Section 5 describes an application to the budget allocation problem.

2 Preliminaries

For an integer $n \in \mathbb{N}$, let $[n]$ denote the set $\{1, \dots, n\}$. In this paper, the symbol E always denotes a (finite) domain of a function.

For a function $w : E \to \mathbb{R}$ and a subset $S \subseteq E$, we define $w(S) = \sum_{e \in S} w(e)$. Similarly, for a vector $\boldsymbol{x} \in \mathbb{R}^E$ and a set $S \subseteq E$, we define $\boldsymbol{x}(S) = \sum_{e \in S} \boldsymbol{x}(e)$. For an element $e \in E$, we define $\mathbf{1}_e$ as the unit vector whose e-th element is 1. For a set $S \subseteq E$, we define $\mathbf{1}_S$ as $\sum_{e \in S} \mathbf{1}_e$.

Let $f : 2^E \to \mathbb{R}$ be a function. For an element $e \in E$, we simply write $f(e)$ to denote $f(\{e\})$. For a set $S \subseteq E$, we define a function $f_S : 2^E \to \mathbb{R}$ as $f_S(T) = f(S \cup T) - f(S)$. We say that f is *submodular* if, for any $S, T \subseteq E$,

$$f(S) + f(T) \geq f(S \cap T) + f(S \cup T).$$

An equivalent condition is the *diminishing return property*, which requires $f_S(e) \geq f_T(e)$ for any $S \subseteq T \subsetneq E$ and $e \in E \setminus T$. We say that f is *linear* if $f(S) = \sum_{e \in S} f(e)$ holds for every $S \subseteq E$. Note that, if f is submodular (resp., linear), then f_S is also submodular (resp., linear).

For a vector $\boldsymbol{x} \in [0,1]^E$, let $R(\boldsymbol{x})$ denote a random set, where each element $e \in E$ is included in the set with probability $\boldsymbol{x}(e)$. For a submodular function $f : 2^E \to \mathbb{R}$, the *multilinear extension* $F : [0,1]^E \to \mathbb{R}$ of f is defined as

$$F(\boldsymbol{x}) := \mathbf{E}[f(R(\boldsymbol{x}))] = \sum_{S \subseteq E} f(S) \prod_{e \in S} \boldsymbol{x}(e) \prod_{e \in E \setminus S} (1 - \boldsymbol{x}(e)).$$

For an element $e \in E$ and a vector $\boldsymbol{x} \in [0,1]^E$, we define $\partial_e F(\boldsymbol{x})$ as the slope of F at $\boldsymbol{x}$ in the direction of $\mathbf{1}_e$. The following fact is well known (see, e.g., [6]):

$$\partial_e F(\boldsymbol{x}) = \frac{F(\boldsymbol{x} \vee \mathbf{1}_e) - F(\boldsymbol{x})}{1 - x_e} = \frac{\mathbf{E}[f_{R(\boldsymbol{x})}(e)]}{1 - x_e}. \tag{1}$$

The following lemma bounds the marginal gain of F when adding $\epsilon \boldsymbol{y}$ to $\boldsymbol{x}$:

Lemma 2.1. *Let $f : 2^E \to \mathbb{R}_+$ be a monotone submodular function and $\boldsymbol{x}, \boldsymbol{y} \in [0,1]^E$ be vectors such that $\boldsymbol{x} + \epsilon \boldsymbol{y} \in [0,1]^E$. Then,*

$$F(\boldsymbol{x} + \epsilon \boldsymbol{y}) - F(\boldsymbol{x}) \geq \epsilon \sum_{e \in E} \boldsymbol{y}(e)\, \mathbf{E}[f_{R(\boldsymbol{x} + \epsilon \boldsymbol{y})}(e)].$$

Proof. Let $e_1, \dots, e_n$ be an arbitrary ordering of elements in E. For $i \in \{0, 1, \dots, n\}$, let $\boldsymbol{x}^i = \boldsymbol{x} + \sum_{j=1}^{i} \epsilon \boldsymbol{y}(e) \mathbf{1}_e$. Note that $\boldsymbol{x}^0 = \boldsymbol{x}$ and $\boldsymbol{x}^n = \boldsymbol{x} + \epsilon \boldsymbol{y}$. Then, we have

$$\begin{aligned}
F(\boldsymbol{x} + \epsilon \boldsymbol{y}) - F(\boldsymbol{x}) &= \sum_{i \in [n]} \partial_{e_i} F(\boldsymbol{x}^{i-1}) \epsilon \boldsymbol{y}(e_i) && \text{(by multilinearity of } F\text{)} \\
&\geq \sum_{i \in [n]} \mathbf{E}[f_{R(\boldsymbol{x}^{i-1})}(e)] \epsilon \boldsymbol{y}(e_i) && \text{(by (1))} \\
&\geq \epsilon \sum_{i \in [n]} \boldsymbol{y}(e_i)\, \mathbf{E}[f_{R(\boldsymbol{x}^n)}(e)]. && \text{(by submodularity of } f\text{)} \\
&= \epsilon \sum_{i \in [n]} \boldsymbol{y}(e_i)\, \mathbf{E}[f_{R(\boldsymbol{x} + \epsilon \boldsymbol{y})}(e)]. && \square
\end{aligned}$$

We frequently use the following form of Chernoff's bound.

Lemma 2.2 (Relative+Additive Chernoff's bound [2])**.** *Let $X_1, \dots, X_n$ be independent random variables such that $X_i \in [0, 1]$ for every $i \in [n]$. Let $X = \frac{1}{n} \sum_{i \in [n]} X_i$ and $\mu = \mathbf{E}[X]$. Then, for any $\alpha \in (0, 1)$ and $\beta > 0$, we have*

$$\Pr[|X - \mu| > \alpha\mu + \beta] \leq 2 \exp\Bigl(-\frac{n\alpha\beta}{3}\Bigr).$$

This immediately gives the following sampling algorithm:

Corollary 2.3. *Suppose that we can obtain independent samples of a random variable X bounded in $[0, d]$. Let $\mu = \mathbf{E}[X]$. Then, there exists an algorithm, denoted by $\textsc{Estimate}_{\alpha,\beta,\delta}(X)$, that, given $\alpha, \beta, \delta \in (0, 1)$, outputs a value $\hat{\mu}$ such that $|\hat{\mu} - \mu| \leq \alpha\mu + \beta d$ with probability at least $1 - \delta$. The number of samples used by the algorithm is $O(\log(1/\delta)/(\alpha\beta))$.*

3 Reduction

In this section, we prove Theorem 1.1 using the following theorem, which gives a joint approximation of a monotone submodular function and a monotone linear function.

Theorem 3.1. *There is an algorithm that, given a monotone submodular function $g : 2^E \to \mathbb{R}_+$, a monotone linear function $\ell : 2^E \to \mathbb{R}_+$, a weight function $w : E \to [0, 1]$, and $\epsilon \in (0, 1)$, outputs a (random) set $S \subseteq E$ with $w(S) \leq 1$ satisfying*

$$\mathbf{E}[g(S) + \ell(S)] \geq \Bigl(1 - \frac{1}{e}\Bigr) g(O) + \ell(O) - \epsilon\bigl(g(O) + \ell(O)\bigr).$$

Here, $O = \arg\max_{T \subseteq E : w(T) \leq 1} \bigl(g(T) + \ell(T)\bigr)$ is an optimal solution. The running time is $O\Bigl(\frac{n^4 \mathrm{polylog}(n)}{(1 - c_g)^2}\Bigr) \cdot \Bigl(\frac{1}{\epsilon} \log \frac{1}{1 - c_g}\Bigr)^{\mathrm{poly}(1/\epsilon)/(1 - c_g)}$, where $n = |E|$.

The proof of Theorem 3.1 is given in Section 4. In the remainder of this section, we prove Theorem 1.1 using Theorem 3.1. The argument is similar to that used in [16], but it is more subtle here because the running time in Theorem 3.1 depends on the curvature of g. We use the following lemma.

Lemma 3.2 (Lemma 2.1 of [16]). *If $f : 2^E \to \mathbb{R}_+$ is a monotone submodular function, then $\sum_{e\in E} f_{E-e}(e) \geq (1-c_f)f(S)$ for all $S \subseteq E$.*

Theorem 3.3. *There is an algorithm that, given a monotone submodular function $f : 2^E \to \mathbb{R}_+$, a weight function $w : E \to [0,1]$, and $\epsilon \in (0,1)$, outputs a (random) set $S \subseteq E$ with $w(S) \leq 1$ satisfying*

$$\mathbf{E}[f(S)] \geq \Big(1 - \frac{c_f}{e} - \epsilon\Big) f(O).$$

Here, $O = \arg\max_{T\subseteq E: w(T)\leq 1} f(T)$ is an optimal solution. The running time is $O\Big(\frac{n^4 \mathrm{polylog}(n)}{(1-c_f)^2}\Big) \cdot \Big(\frac{1}{\epsilon}\log\frac{1}{1-c_f}\Big)^{\mathrm{poly}(1/\epsilon)/(1-c_f)}$, where $n = |E|$.

Proof. Define the functions $g, \ell : 2^E \to \mathbb{R}_+$ such that

$$\ell(S) = \Big(1 - \frac{\epsilon}{2}\Big) \sum_{e\in S} f_{E-e}(e) \quad \text{and} \quad g(S) = f(S) - \ell(S)$$

for every $S \subseteq E$.

It is not hard to see that ℓ is a nonnegative monotone linear function and that g is a nonnegative monotone submodular function. Moreover, the curvature of g is

$$c_g = 1 - \min_{e\in E} \frac{g_{E-e}(e)}{g(e)} = 1 - \min_{e\in E} \frac{f_{E-e}(e) - (1-\epsilon/2)f_{E-e}(e)}{f(e) - (1-\epsilon/2)f_{E-e}(e)} \leq 1 - \frac{\epsilon}{2}\min_{e\in E}\frac{f_{E-e}(e)}{f(e)} = 1 - \frac{\epsilon(1-c_f)}{2}.$$

Further, Lemma 3.2 implies that for any set $S \subseteq E$,

$$\ell(S) = \Big(1 - \frac{\epsilon}{2}\Big)\sum_{e\in S} f_{E-e}(e) \geq \Big(1 - \frac{\epsilon}{2}\Big)(1-c_f)f(S) \geq \Big(1 - c_f - \frac{\epsilon}{2}\Big) f(S).$$

By applying Theorem 3.1 to g, ℓ, w, and $\epsilon/2$, we can find a (random) set $S \subseteq E$ with $w(S) \leq 1$ satisfying

$$\begin{aligned}
\mathbf{E}[f(S)] = \mathbf{E}[g(S) + \ell(S)] &\geq \Big(1 - \frac{1}{e}\Big) g(O) + \ell(O) - \frac{\epsilon}{2}\big(g(O) + \ell(O)\big) \\
&= \Big(1 - \frac{1}{e}\Big) f(O) + \frac{1}{e}\ell(O) - \frac{\epsilon}{2} f(O) \\
&\geq \Big(1 - \frac{1}{e}\Big) f(O) + \frac{1 - c_f - \epsilon/2}{e} f(O) - \frac{\epsilon}{2} f(O) \\
&\geq \Big(1 - \frac{c_f}{e} - \epsilon\Big) f(O).
\end{aligned}$$

The running time is clearly as stated. ∎

Now, we prove our main theorem.

Proof of Theorem 1.1. If $c_f < 1 - e\epsilon$, then we run the algorithm in Theorem 3.3. The approximation factor is $1 - c_f/e - \epsilon$ and the running time is $O\Big(\frac{n^4\mathrm{polylog}(n)}{(1-c_f)^2}\Big) \cdot \Big(\frac{1}{\epsilon}\log\frac{1}{1-c_f}\Big)^{\mathrm{poly}(1/\epsilon)/(1-c_f)} = O\Big(n^4\mathrm{polylog}(n) \cdot \big(\frac{1}{\epsilon}\big)^{\mathrm{poly}(1/\epsilon)}\Big)$.

If $c_f \geq 1 - e\epsilon$, then we simply run the $O(n^5)$-time $(1-1/e)$-approximation algorithm presented in [15]. Then, the approximation factor is

$$1 - \frac{1}{e} \geq 1 - \frac{1 - e\epsilon}{e} - \epsilon \geq 1 - \frac{c_f}{e} - \epsilon.$$

In both cases, the approximation factor is at least $1 - c_f/e - \epsilon$ whereas the running time is as desired. □

4 Proof of Theorem 3.1

In this section, we prove Theorem 3.1. Throughout this section, $G : [0,1]^E \to \mathbb{R}_+$ denotes the multilinear extension of g, while $L : [0,1]^E \to \mathbb{R}_+$ and $W : [0,1]^E \to \mathbb{R}_+$ denote the following linear functions:

$$L(\boldsymbol{x}) = \sum_{e \in E} \boldsymbol{x}(e)\ell(e) \quad \text{and} \quad W(\boldsymbol{x}) = \sum_{e \in E} \boldsymbol{x}(e)w(e).$$

Furthermore, we define $d_g = \max_{e \in E} g(e)$, $d_\ell = \max_{e \in E} \ell(e)$, and $d_{g,\ell} = \max(d_g, d_\ell)$. Note that we have $g(O) \leq nd_g \leq nd_{g,\ell}$, $\ell(O) \leq nd_\ell \leq nd_{g,\ell}$, and $d_{g,\ell} \leq g(O) + \ell(O) \leq 2nd_{g,\ell}$. Recall that $n = |E|$.

In Section 4.1, we argue that we need to deal with "small" and "large" elements separately in order to guarantee that we get a set satisfying the knapsack constraint after rounding. Our algorithm updates a vector $\boldsymbol{x} \in [0,1]^E$ in several iterations and then rounds it. In Section 4.2, we present an algorithm that updates $\boldsymbol{x}$ by adding a vector supported on small elements. In Section 4.3, we present the entire algorithm that computes the vector $\boldsymbol{x}$ by taking large elements into account. Then, in Section 4.4, we describe our rounding procedure. We need to guess several parameters when running the algorithms in Section 4.2 and 4.3. Our final algorithm with the guessing process is presented in Section 4.5.

4.1 Small elements

Our algorithm computes a vector $\boldsymbol{x} \in [0,1]^E$ with $W(\boldsymbol{x}) \leq 1$ and then rounds it to a set. A natural rounding method is to simply output the random set $R(\boldsymbol{x})$. Then, we can guarantee that the expected objective values $\mathbf{E}[g(R(\boldsymbol{x}))]$ and $\mathbf{E}[\ell(R(\boldsymbol{x}))]$ are sufficiently large and the expected weight $\mathbf{E}[w(R(\boldsymbol{x}))]$ is at most one. However, we cannot guarantee the concentration of $w(R(\boldsymbol{x}))$ because some elements have large contributions to the weight. To resolve this issue, we say that elements in E are *small* if

$$g(e) \leq \epsilon^6 g(O) \quad \text{and} \quad \ell(e) \leq \epsilon^6 \ell(O).$$

Then, we can freely remove some of small elements for decreasing the weight without decreasing the value significantly. Further, we can prove that the number of large elements is bounded by a polynomial in ϵ and c_g.

An issue here is that we do not know O; hence, we cannot determine whether an element is small. To resolve this issue, we guess the values of $g(O)$ and $\ell(O)$. Without loss of generality, we can assume that $\epsilon/n = (1-\epsilon)^k$ for some integer k; otherwise, we slightly decrease the value of ϵ. Then, we define a set $V_{\epsilon,n}(g,\ell) = \{nd_{g,\ell}, (1-\epsilon)nd, \ldots, \epsilon d_{g,\ell}, 0\}$, and we use the values in $V_\epsilon(g,\ell)$ to

guess $g(O)$ and $\ell(O)$. Since $g(O) \leq nd_{g,\ell}$ and $\ell(O) \leq nd_{g,\ell}$ hold, there exist some $v_g, v_\ell \in V_{\epsilon,n}(g,\ell)$ such that

$$(1-\epsilon)v_g - \epsilon d_{g,\ell} \leq g(O) \leq v_g \quad \text{and} \quad (1-\epsilon)v_\ell - \epsilon d_{g,\ell} \leq \ell(O) \leq v_\ell. \tag{2}$$

We say that an element $e \in E$ is *small with respect to* (v_g, v_ℓ) if

$$g(e) \leq \epsilon^6 v_g \quad \text{and} \quad \ell(e) \leq \epsilon^6 v_\ell.$$

Otherwise, we say that an element $e \in E$ is *large with respect to* (v_g, v_ℓ). Let $E_{\mathrm{L}}(v_g, v_\ell) \subseteq E$ and $E_{\mathrm{S}}(v_g, v_\ell) \subseteq E$ be the sets of large and small elements, respectively, with respect to (v_g, v_ℓ). Further, we define $O_{\mathrm{L}}(v_g, v_\ell) = E_{\mathrm{L}}(v_g, v_\ell) \cap O$ and $O_{\mathrm{S}}(v_g, v_\ell) = E_{\mathrm{S}}(v_g, v_\ell) \cap O$. We omit v_g and v_ℓ from these notations if they are clear from the context.

When v_g and v_ℓ satisfy (2), we can upper bound the number of large elements in O:

Lemma 4.1. *If v_g and v_ℓ satisfy (2), then we have $|O_{\mathrm{L}}| = O\big(\frac{1}{(1-c_g)\epsilon^6}\big)$.*

Proof. Let m_ℓ be the number of elements $e \in O$ with $\ell(e) > \epsilon^6 v_\ell$. Then, we have $\epsilon^6 v_\ell m_\ell \leq \ell(O)$. Since $v_\ell \geq \ell(O)$, we have $m_\ell \leq 1/\epsilon^6$.

Let $\{o_1, \ldots, o_{m_g}\}$ be the set of elements $e \in O$ with $g(e) > \epsilon^6 v_g$. Then, we have

$$g(O) \geq \sum_{i \in [m_g]} g_{\{o_1,\ldots,o_{i-1}\}}(o_i) \geq (1-c_g) \sum_{i \in [m_g]} g(o_i) \geq (1-c_g)\epsilon^6 v_g m_g.$$

Since $v_g \geq g(O)$, we have $m_g \leq \frac{1}{(1-c_g)\epsilon^6}$.

Then, we have $|O_{\mathrm{L}}| \leq m_g + m_\ell = O(\frac{1}{(1-c_g)\epsilon^6})$. □

In addition to the values of $g(O)$ and $\ell(O)$, the value of $|O_{\mathrm{L}}|$ is not also not known. However, we can easily guess it because there are only $O(\frac{1}{(1-c_g)\epsilon^6})$ choices from Lemma 4.1. We use the symbol m to denote the guessed value of $|O_{\mathrm{L}}|$.

For each choice of v_g, v_ℓ, and m, we compute a (random) set that jointly maximizes g and ℓ, and the final output is the best one among them. Since $|V_{\epsilon,n}(g,\ell)| = O(\log_{1/(1-\epsilon)}(n/\epsilon)) = O(\log(n/\epsilon)/\epsilon)$, this guessing process makes the running time $O\Big(\big(\log(n/\epsilon)/\epsilon\big)^2 \cdot \frac{1}{(1-c_g)\epsilon^6}\Big) = O\Big(\frac{\log^2(n/\epsilon)}{(1-c_g)\epsilon^8}\Big)$ times larger. The details will be explained in Section 4.5.

4.2 Subroutine for handling small elements

Here, we explain a subroutine that finds a vector $\boldsymbol{v} \in [0,1]^E$ supported on the set E_S of small elements (with respect to the current guesses v_g and v_ℓ) in order to update the current vector $\boldsymbol{x} \in \mathbb{R}^E$. We want $\boldsymbol{v}$ to satisfy the following properties: (i) $\sum_{e \in E_{\mathrm{S}}} \boldsymbol{v}(e)\, \mathbf{E}[g_{R(\boldsymbol{x})}(e)] \geq \mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})]$, (ii) $L(\boldsymbol{v}) \geq \ell(O_{\mathrm{S}})$, and (iii) $W(\boldsymbol{v}) \leq w(O_{\mathrm{S}})$.

There are several issues in finding such a vector $\boldsymbol{v}$: We cannot exactly calculate $\mathbf{E}[g_{R(\boldsymbol{x})}(e)]$ ($e \in E_{\mathrm{S}}$); hence, we need to estimate it. Further, we do not know the values $\mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})]$ and $\ell(O_{\mathrm{S}})$. In the subroutine presented here, we assume that their guessed values, denoted by γ and λ, respectively, are given as a part of the input. Once we succeed in accurately estimating $\mathbf{E}[g_{R(\boldsymbol{x})}(e)]$ ($e \in E_{\mathrm{S}}$) and the given guessed values γ and λ are sufficiently accurate, we can find the desired vector $\boldsymbol{v}$ by solving a linear program. A detailed description of the subroutine is given in Algorithm 1.

Now, we analyze Algorithm 1. From Corollary 2.3 and the fact that $g_S(e) \leq d_g \leq d_{g,\ell}$ for every $S \subseteq E$ and $e \in E$, we have the following:

Algorithm 1 SmallElements$_{\epsilon,\delta}(g, \ell, w, E_{\mathrm{S}}, \gamma, \lambda, \boldsymbol{x})$

Input: A monotone submodular function $g : 2^E \to \mathbb{R}_+$, a monotone linear function $\ell : 2^E \to \mathbb{R}$, a weight function $w : E \to [0, 1]$, $\epsilon, \delta \in (0, 1)$, a set of small elements E_{S}, guessed values γ, λ, and a vector $\boldsymbol{x} \in [0, 1]^E$.

Output: A vector $\boldsymbol{v} \in [0, 1]^E$.

1: For each $e \in E_{\mathrm{S}}$, let $\boldsymbol{\theta}(e) = \textsc{Estimate}_{\epsilon,\epsilon/n,\delta/n}(g_{R(\boldsymbol{x})}(e))$.

2: Find a vector $\boldsymbol{v} \in [0, 1]^E$ supported on E_{S} that minimizes $W(\boldsymbol{v})$ subject to

$$\boldsymbol{v} \cdot \boldsymbol{\theta} \geq (1 - \epsilon)\gamma - \epsilon d_{g,\ell} \quad \text{and} \quad L(\boldsymbol{v}) \geq \lambda.$$

by linear programming.

3: **return** $\boldsymbol{v}$.

Proposition 4.2. *With probability at least $1 - \delta$, we have*

$$(1 - \epsilon)\,\mathbf{E}[g_{R(\boldsymbol{x})}(e)] - \frac{\epsilon d_{g,\ell}}{n} \leq \boldsymbol{\theta}(e) \leq (1 + \epsilon)\,\mathbf{E}[g_{R(\boldsymbol{x})}(e)] + \frac{\epsilon d_{g,\ell}}{n}$$

for every $e \in E_{\mathrm{S}}$.

We formalize the concept that γ and λ are sufficiently accurate, and then show that Algorithm 1 outputs a desired vector with accurate γ and λ.

Definition 4.3. *We say that γ and λ are* good guesses *if*

$$\mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})] \geq \gamma \geq (1 - \epsilon)\,\mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})] - \epsilon d_{g,\ell} \quad \textit{and} \quad \ell(O_{\mathrm{S}}) \geq \lambda \geq (1 - \epsilon)\ell(O_{\mathrm{S}}) - \epsilon d_{g,\ell}$$

hold, respectively.

Since $\mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})] \leq n d_{g,\ell}$ and $\ell(O_{\mathrm{S}}) \leq n d_{g,\ell}$ hold, we can find good guesses by trying all the values in the set $V_{\epsilon,n}(g, \ell)$.

Lemma 4.4. *Suppose that γ and λ are good guesses. Then, Algorithm 1 returns a vector $\boldsymbol{v} \in [0, 1]^E$ supported on E_{S} such that*

(i) $\sum_{e \in E_{\mathrm{S}}} \boldsymbol{v}(e)\,\mathbf{E}[g_{R(\boldsymbol{x})}(e)] \geq (1 - \epsilon)^3\,\mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})] - 3\epsilon d_{g,\ell}$,

(ii) $L(\boldsymbol{v}) \geq (1 - \epsilon)\ell(O_{\mathrm{S}}) - \epsilon d_{g,\ell}$, *and*

(iii) $W(\boldsymbol{v}) \leq w(O_{\mathrm{S}})$,

with probability at least $1 - \delta$. The time complexity of Algorithm 1 is $O(n^4 + n^2 \log(n/\delta)/\epsilon^2)$.

Proof. With probability at least $1 - \delta$, the consequence of Proposition 4.2 holds. In what follows, we assume that this occurs.

The vector $\mathbf{1}_{O_{\mathrm{S}}}$ satisfies

$$\begin{aligned}\mathbf{1}_{O_{\mathrm{S}}} \cdot \boldsymbol{\theta} = \sum_{e \in O_{\mathrm{S}}} \boldsymbol{\theta}(e) &\geq \sum_{e \in O_{\mathrm{S}}} \left((1 - \epsilon)\,\mathbf{E}[g_{R(\boldsymbol{x})}(e)] - \frac{\epsilon d_{g,\ell}}{n}\right) \\ &\geq (1 - \epsilon)\,\mathbf{E}[g_{R(\boldsymbol{x})}(O_{\mathrm{S}})] - \epsilon d_{g,\ell} \geq (1 - \epsilon)\gamma - \epsilon d_{g,\ell}.\end{aligned}$$

Furthermore, we have $L(\mathbf{1}_{O_S}) = \ell(O_S) \geq \lambda$. Hence, the vector $\boldsymbol{v}$ is well defined, and in particular, we have $W(\boldsymbol{v}) \leq W(O_S)$.

Then, we have

$$\begin{aligned}\sum_{e\in E_S} \boldsymbol{v}(e)\,\mathbf{E}[g_{R(\boldsymbol{x})}(e)] &\geq (1-\epsilon)\sum_{e\in E_S} \boldsymbol{v}(e)\Big(\boldsymbol{\theta}(e) - \frac{\epsilon d_{g,\ell}}{n}\Big) \geq (1-\epsilon)\sum_{e\in E_S} \boldsymbol{v}(e)\boldsymbol{\theta}(e) - \epsilon d_{g,\ell}\\ &\geq (1-\epsilon)\big((1-\epsilon)\gamma - \epsilon d_{g,\ell}\big) - \epsilon d_{g,\ell} \geq (1-\epsilon)^2\gamma - 2\epsilon d_{g,\ell}\\ &\geq (1-\epsilon)^2\big((1-\epsilon)\,\mathbf{E}[g_{R(\boldsymbol{x})}(O_S)] - \epsilon d_{g,\ell}\big) - 2\epsilon d_{g,\ell}\\ &\geq (1-\epsilon)^3\,\mathbf{E}[g_{R(\boldsymbol{x})}(O_S)] - 3\epsilon d_{g,\ell}.\end{aligned}$$

It is easy to confirm (ii) and (iii). The time complexity for computing $\boldsymbol{\theta}$ is $O(n^2\log(n/\delta)/\epsilon^2)$ from Corollary 2.3, and the time complexity for solving the linear program is $O(n^4)$ by using the ellipsoid method. The total time complexity is bounded by $O(n^4 + n^2\log(n/\delta)/\epsilon^2)$. □

4.3 Continuous greedy algorithm with guessing

In this section, we present an algorithm whose goal is to output a vector $\boldsymbol{x} \in [0,1]^E$ such that (i) $G(\boldsymbol{x}) \geq (1-1/e)g(O)$, (ii) $L(\boldsymbol{x}) \geq \ell(O)$, and (iii) $W(\boldsymbol{x}) \leq w(O)$.

Our algorithm is a variant of the continuous greedy algorithm [3] but differs in the following aspects: we consider two functions g and ℓ, and we handle large and small elements separately. Let m be an integer given as a parameter, which is a guessed value of $|O_L|$ (with respect to the current values of v_g and v_ℓ). Then, we make copies $E_1,\dots,E_m$ of E and define a set $\widehat{E} = \bigcup_{i\in[m]} E_i \cup E_S$. Then, we define a function $\widehat{g} : 2^{\widehat{E}} \to \mathbb{R}_+$ as $\widehat{g}(S_1,S_2,\dots,S_m,S_S) = g(S_1\cup\cdots\cup S_m\cup S_S)$. We note that $\widehat{g}$ is a monotone submodular function. Let $\widehat{G}$ be the multilinear extension of $\widehat{g}$.

We introduce a vector $\boldsymbol{y}_i \in [0,1]^E$ for each $i\in[m]$ and another vector $\boldsymbol{z}\in[0,1]^E$. We always guarantee that $\boldsymbol{y}_i$ ($i\in[m]$) is supported on E_i and $\boldsymbol{z}$ is supported on E_S. Our algorithm runs in $1/\epsilon$ iterations and updates the vectors $\boldsymbol{y}_i$ ($i\in[m]$) and $\boldsymbol{z}$ in each iteration. Here, we assume that $1/\epsilon$ is an integer; otherwise, we slightly decrease ϵ. The final output is the sequence of vectors $(\boldsymbol{y}_1,\dots,\boldsymbol{y}_m,\boldsymbol{z})$, and their sum $\boldsymbol{x} := \sum_{i\in[m]}\boldsymbol{y}_i + \boldsymbol{z}$ will satisfy the conditions stated initially in this section. We call the first iteration the iteration at time 0, the second iteration the iteration at time ϵ, and so on.

To explain how we update the vectors, we introduce several notations. For $t\in\{0,\epsilon,\dots,1\}$, we define $\boldsymbol{y}_i^t$ ($i\in[m]$) and $\boldsymbol{z}^t$ as the vectors $\boldsymbol{y}_i$ and $\boldsymbol{z}$ immediately before the iteration at time t. We note that $\boldsymbol{y}_i^0 = \mathbf{0}$ ($i\in[m]$) and $\boldsymbol{z}^0 = \mathbf{0}$ hold. We define $\boldsymbol{y}_i^1$ ($i\in[m]$) and $\boldsymbol{z}^1$ as $\boldsymbol{y}_i$ ($i\in[m]$) and $\boldsymbol{z}$, respectively, after the iteration at time $1-\epsilon$. Note that the algorithm outputs the sequence of vectors $(\boldsymbol{y}_1^1,\dots,\boldsymbol{y}_m^1,\boldsymbol{z}^1)$. Then, we define $\boldsymbol{x}^t = \sum_{i\in[m]}\boldsymbol{y}_i^t + \boldsymbol{z}^t$. Further, for $t\in\{0,\epsilon,\dots,1-\epsilon\}$ and $i\in\{0,1,\dots,m\}$, we define $\boldsymbol{x}_i^t = \sum_{j\leq i}\boldsymbol{y}_j^{t+\epsilon} + \sum_{j>i}\boldsymbol{y}_j^t + \boldsymbol{z}^t$, i.e., the vector obtained after the iteration at time $t-\epsilon$ followed by updating $\boldsymbol{y}_1,\dots,\boldsymbol{y}_i$. Note that $\boldsymbol{x}_0^t = \boldsymbol{x}^t$.

As in the argument in Section 4.1, we try all possible values for guessing $|O_L|$. Hence, in what follows, we assume that the guessed value m is correct, i.e., $m = |O_L|$.

Let $o_1,\dots,o_m$ be the large elements in O, i.e., $O_L = \{o_1,\dots,o_m\}$. For $i\in[m]$, let $\widehat{o}_i$ be the copy of o_i in E_i. Then, we define $\widehat{O}_L = \{\widehat{o}_1,\dots,\widehat{o}_m\}$ and $\widehat{O} = \widehat{O}_L\cup O_S \subseteq \widehat{E}$. For each $i\in[m]$, we update the vector $\boldsymbol{y}_i^t$ to $\boldsymbol{y}_i^{t+\epsilon}$ by finding an element $e_i^t\in E_i$ and adding the vector $\epsilon\mathbf{1}_{e_i^t}$. Here, we want the element e_i^t to satisfy (i) $\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(e_i^t)] \geq \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(\widehat{o}_i)]$, (ii) $\ell(e_i^t)\geq\ell(o_i)$, and (iii)

Algorithm 2 GuessingContinuousGreedy$_{\epsilon,\delta}(g, \ell, w, E_{\mathrm{L}}, E_{\mathrm{S}}, m, \{\gamma_i^t\}, \{\gamma_{\mathrm{S}}^t\}, \{\lambda_i\}, \lambda_{\mathrm{S}})$

Input: A monotone submodular function $g : 2^E \to \mathbb{R}_+$, a monotone linear function $\ell : 2^E \to \mathbb{R}$, a weight function $w : E \to [0, 1]$, $\epsilon, \delta \in (0, 1)$, a set of large and small elements E_{L} and E_{S}, an integer $m \in \mathbb{N}$, guessed values $\{\gamma_i^t\}_{i\in[m],t\in\{0,\epsilon,\dots,1-\epsilon\}}$, $\{\gamma_{\mathrm{S}}^t\}_{t\in\{0,\epsilon,\dots,1-\epsilon\}}$, $\{\lambda_i\}_{i\in[m]}$, and λ_{S}.

Output: A vector $\boldsymbol{x} \in [0, 1]^E$.

$\boldsymbol{y}_i \leftarrow \mathbf{0} \in [0, 1]^{\widehat{E}}$ for $i \in [m]$ and $\boldsymbol{z} \leftarrow \mathbf{0} \in [0, 1]^{\widehat{E}}$.
for $(t \leftarrow 0; t \le 1 - \epsilon; t \leftarrow t + \epsilon)$ **do**
  for $(i \leftarrow 1; i \le m; i \leftarrow i + 1)$ **do**
    $\boldsymbol{\theta}_i(e) \leftarrow \textsc{Estimate}_{\epsilon,\epsilon/m,\epsilon\delta/(2nm)}(\mathbf{E}[\widehat{g}_{R(\boldsymbol{x})}](e))$ for each $e \in E_i$.
    Let $e = \arg\min\{w(e) \mid e \in E_i, \boldsymbol{\theta}_i(e) \ge (1 - \epsilon)\gamma_i^t - \epsilon d/m, \ell(e) \ge \lambda_i\}$.
    $\boldsymbol{y}_i \leftarrow \boldsymbol{y}_i + \epsilon \mathbf{1}_e$.
  $\boldsymbol{v} \leftarrow \textsc{SmallElements}_{\epsilon,\epsilon\delta/2}(\widehat{g}, \ell, w, E_{\mathrm{S}}, \gamma_{\mathrm{S}}^t, \lambda_{\mathrm{S}}, \sum_{i\in[m]} \boldsymbol{y}_i + \boldsymbol{z})$.
  $\boldsymbol{z} \leftarrow \boldsymbol{z} + \epsilon \boldsymbol{v}$.
return $(\boldsymbol{y}_1, \dots, \boldsymbol{y}_m, \boldsymbol{z})$.

$w(e_i^t) \le w(o_i)$. As we do not know the values of $\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(\widehat{o}_i)]$ and $\ell(o_i)$, the algorithm requires their guessed values γ_i^t and λ_i, respectively. We do not have to guess $w(o_i)$ because we will choose the element with the minimum weight satisfying (i) and (ii).

Then, we update the vector $\boldsymbol{z}^t$ to $\boldsymbol{z}^{t+\epsilon}$ by finding a vector $\boldsymbol{v}^t$ and adding the vector $\epsilon\boldsymbol{v}^t$. Here, we want the vector $\boldsymbol{v}$ to satisfy (i) $\sum_{e\in E_{\mathrm{S}}} \boldsymbol{v}(e)\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x})}(e)] \ge \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(O_S)]$ (note that $O_S \subseteq E_S \subseteq \widehat{E}$), (ii) $L(\boldsymbol{v}) \ge \ell(O_{\mathrm{S}})$, and (iii) $W(\boldsymbol{v}) \le w(O_{\mathrm{S}})$. Such a vector can be found by calling SmallElements with the guessed values γ_{S}^t and λ_{S} for $\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(O_S)]$ and $\ell(O_{\mathrm{S}})$, respectively. A detailed description of the algorithm is given in Algorithm 2.

We will show that $\widehat{G}(\boldsymbol{x}^{t+\epsilon}) - G(\boldsymbol{x}^t) \ge \epsilon(g(O) - \widehat{G}(\boldsymbol{x}^{t+\epsilon}))$, which is sufficient to show that $\widehat{G}(\boldsymbol{x}^1)$ is close to the $(1 - 1/e)$-approximation to $g(O)$. We will also show that $L(\boldsymbol{x}^1) \ge \ell(O)$ and $W(\boldsymbol{x}) \le w(O)$.

For $t \in \{0, \epsilon, \dots, 1 - \epsilon\}$ and $i \in [m]$, let $\boldsymbol{\theta}_i^t$ be the $\boldsymbol{\theta}_i$ used in the iteration at time t. From Lemma 2.2 and the union bound, we immediately have the following:

Proposition 4.5. *With probability at least $1 - \delta/2$, we have*

$$(1 - \epsilon)\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(e)] - \frac{\epsilon d_{g,\ell}}{m} \le \boldsymbol{\theta}_i^t(e) \le (1 + \epsilon)\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(e)] + \frac{\epsilon d_{g,\ell}}{m}$$

for every $t \in \{0, \epsilon, \dots, 1 - \epsilon\}$, $i \in [m]$, and $e \in E_i$.

We formalize the concept that γ_i^t and λ_i are sufficiently accurate.

Definition 4.6. *For $t \in \{0, \epsilon, \dots, 1 - \epsilon\}$ and $i \in [m]$, we say that γ_i^t is a* good guess *if*

$$\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(\widehat{o}_i)] \ge \gamma_i^t \ge (1 - \epsilon)\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{i-1}^t)}(\widehat{o}_i)] - \frac{\epsilon d_{g,\ell}}{m}$$

holds. For $i \in [m]$, we say that λ_i is a good guess *if*

$$\ell(o_i) \ge \lambda_i \ge (1 - \epsilon)\ell(o_i) - \frac{\epsilon d_{g,\ell}}{m}$$

holds.

Since $\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^t_{i-1})}(o_i)] \le d_{g,\ell}$ and $\ell(o_i) \le d_{g,\ell}$ hold for every $i \in [m]$, we can find good guesses by trying all the values in the set $V_{\epsilon,m}(g,\ell)/m := \{v/m \mid v \in V_{\epsilon,m}(g,\ell)\}$.

Lemma 4.7. *Suppose that the consequence of Proposition 4.5 holds and that $\{\gamma^t_i\}$ and $\{\lambda_i\}$ are good guesses. Then, for every $t \in \{0,\epsilon,\dots,1-\epsilon\}$, we have the following:*

(i) $\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^t_{i-1})}(e^t_i)] \ge (1-\epsilon)^3\,\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^t_{i-1})}(\widehat{o}_i)] - 3\epsilon d_{g,\ell}/m$,

(ii) $\ell(e^t_i) \ge (1-\epsilon)\ell(o_i) - \epsilon d_{g,\ell}/m$ *for* $i \in [m]$, *and*

(iii) $w(e^t_i) \le w(o_i)$ *for* $i \in [m]$.

Proof. Fix $t \in \{0,\epsilon,\dots,1-\epsilon\}$ and $i \in [m]$. Note that we have

$$\boldsymbol{\theta}^t_i(o_i) \ge (1-\epsilon)\,\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^t_{i-1})}(\widehat{o}_i)] - \frac{\epsilon d_{g,\ell}}{m} \ge (1-\epsilon)\gamma^t_i - \frac{\epsilon d_{g,\ell}}{m}$$

and $\ell(o_i) \ge \lambda_i$. Since o_i (in E_i) is a candidate for e^t_i, the element e^t_i is well defined. In particular, we have $w(e^t_i) \le w(o_i)$ because e^t_i is chosen as the element with the minimum element satisfying the conditions.

We have

$$(1+\epsilon)\,\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^t_{i-1})}(e^t_i)] + \frac{\epsilon d_{g,\ell}}{m} \ge \boldsymbol{\theta}^t_i(e^t_i) \ge (1-\epsilon)\gamma^t_i - \frac{\epsilon d_{g,\ell}}{m} \ge (1-\epsilon)^2\,\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^t_i)}(o_i)] - \frac{2\epsilon d_{g,\ell}}{m}.$$

Rearranging this inequality, we get (i). Further, (ii) is immediate from the fact that $\ell(e^t_i) \ge \lambda_i$. □

We say that γ^t_{S} for $t \in \{0,\epsilon,\dots,1-\epsilon\}$ and λ_{S} are *good guesses* if they are good guesses in the sense of Definition 4.3. Then, we have the following:

Lemma 4.8. *Suppose that $\{\gamma^t_i\},\{\lambda_i\}$, $\{\gamma^t_{\mathrm{S}}\}$, and λ_{S} are good guesses. Then, Algorithm 2 returns vectors $\boldsymbol{y}_1,\dots,\boldsymbol{y}_m,\boldsymbol{z}$ such that $\boldsymbol{x} := \sum_{i\in[m]}\boldsymbol{y}_i + \boldsymbol{z}$ satisfies the following:*

(i) $\widehat{G}(\boldsymbol{x}) \ge (1-1/e-O(\epsilon))g(O) - 6\epsilon d_{g,\ell}$,

(ii) $L(\boldsymbol{x}) \ge (1-O(\epsilon))\ell(O) - 2\epsilon d_{g,\ell}$, *and*

(iii) $W(\boldsymbol{x}) \le w(O)$,

with probability at least $1-\delta$. The running time is $O(\frac{nm^2}{\epsilon^3}\log\frac{nm}{\epsilon\delta} + \frac{n^4}{\epsilon} + \frac{n^2}{\epsilon^3}\log\frac{1}{\epsilon\delta})$.

Proof. With probability $1-\delta/2$, the consequence of Proposition 4.5 holds. Further, with probability $1-\delta/2$, all the invocations of SmallElements succeed in outputting vectors with the guarantees in Lemma 4.4. By the union bound, all these occur with probability at least $1-\delta$. In what follows, we assume that this occurs.

First, we check (i). For each $t \in \{0, \epsilon, \dots, 1-\epsilon\}$, we have

$$\begin{aligned}
&\widehat{G}(\boldsymbol{x}_m^t) - \widehat{G}(\boldsymbol{x}^t) = \sum_{j=1}^{m} \Big(\widehat{G}(\boldsymbol{x}_{j-1}^t + \epsilon \mathbf{1}_{e_j^t}) - \widehat{G}(\boldsymbol{x}_{j-1}^t) \Big) \\
&\geq \epsilon \sum_{j=1}^{m} \Big(\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{j-1}^t)}(e_j^t)] \Big) && \text{(By concavity of } \widehat{G}) \\
&\geq \epsilon(1-\epsilon)^3 \sum_{j=1}^{m} \Big(\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_{j-1}^t)}(\widehat{o}_j)] - \frac{3\epsilon d_{g,\ell}}{m} \Big) && \text{(By (i) of Lemma 4.7)} \\
&\geq \epsilon(1-\epsilon)^3 \sum_{j=1}^{m} \Big(\mathbf{E}\big[\widehat{g}_{R(\boldsymbol{x}_m^t) \cup \{\widehat{o}_k \mid k \in [j-1]\}}(\widehat{o}_j)\big] \Big) - 3\epsilon^2 d_{g,\ell} \\
&= \epsilon(1-\epsilon)^3 \Big(\mathbf{E}\big[\widehat{g}(R(\boldsymbol{x}_m^t) \cup \widehat{O}_{\mathrm{L}}) - \widehat{g}(R(\boldsymbol{x}_m^t))\big] \Big) - 3\epsilon^2 d_{g,\ell} \\
&\geq \epsilon(1-\epsilon)^3 (\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_m^t)}(\widehat{O}_{\mathrm{L}})]) - 3\epsilon^2 d_{g,\ell}.
\end{aligned}$$

For each $t \in \{0, \epsilon, \dots, 1-\epsilon\}$, we have

$$\begin{aligned}
&\widehat{G}(\boldsymbol{x}^{t+\epsilon}) - \widehat{G}(\boldsymbol{x}_m^t) = \widehat{G}(\boldsymbol{x}_m^t + \epsilon \boldsymbol{v}) - \widehat{G}(\boldsymbol{x}_m^t) \\
&\geq \epsilon \sum_{e \in E} \boldsymbol{v}(e)\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^{t+\epsilon})}(e)] && \text{(by Lemma 2.1)} \\
&\geq \epsilon \big((1-\epsilon)^3\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^{t+\epsilon})}(O_{\mathrm{S}})] - 3\epsilon d_{g,\ell} \big) && \text{(by (i) of Lemma 4.4)} \\
&\geq \epsilon(1-\epsilon)^3\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^{t+\epsilon})}(O_{\mathrm{S}})] - 3\epsilon^2 d_{g,\ell}.
\end{aligned}$$

Combining these two inequalities, we get

$$\begin{aligned}
&\widehat{G}(\boldsymbol{x}^{t+\epsilon}) - \widehat{G}(\boldsymbol{x}^t) \\
&\geq \epsilon(1-\epsilon)^3 (\mathbf{E}[\widehat{g}_{R(\boldsymbol{x}_m^t)}(\widehat{O}_{\mathrm{L}})]) - 3\epsilon^2 d_{g,\ell} + \epsilon(1-\epsilon)^3\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^{t+\epsilon})}(O_{\mathrm{S}})] - 3\epsilon^2 d_{g,\ell} \\
&\geq \epsilon(1-\epsilon)^3\, \mathbf{E}[\widehat{g}_{R(\boldsymbol{x}^{t+\epsilon})}(\widehat{O})] - 6\epsilon^2 d_{g,\ell} \\
&\geq \epsilon(1-\epsilon)^3 \big(\widehat{g}(\widehat{O}) - \widehat{G}(\boldsymbol{x}^{t+\epsilon}) \big) - 6\epsilon^2 d_{g,\ell} = \epsilon(1-\epsilon)^3 \big(g(O) - \widehat{G}(\boldsymbol{x}^{t+\epsilon}) \big) - 6\epsilon^2 d_{g,\ell}.
\end{aligned}$$

Rewriting the above inequality, we get

$$g(O) - \frac{\beta}{\alpha} - \widehat{G}(\boldsymbol{x}^{t+\epsilon}) \leq \frac{g(O) - \frac{\beta}{\alpha} - \widehat{G}(\boldsymbol{x}^t)}{1+\alpha},$$

where $\alpha = \epsilon(1-\epsilon)^3$ and $\beta = 6\epsilon^2 d_{g,\ell}$. Then, by induction, we can prove that

$$g(O) - \frac{\beta}{\alpha} - \widehat{G}(\boldsymbol{x}^t) \leq \frac{1}{(1+\alpha)^{t/\epsilon}} \Big(g(O) - \frac{\beta}{\alpha} \Big).$$

Substituting $t = 1$ and rewriting once again, we get

$$\widehat{G}(\boldsymbol{x}^1) \geq \Big(1 - \frac{1}{(1+\alpha)^{1/\epsilon}} \Big) (g(O) - \frac{\beta}{\alpha}) \geq \Big(1 - \frac{1}{e} - O(\epsilon) \Big) \Big(g(O) - \frac{\beta}{\alpha} \Big) \geq \Big(1 - \frac{1}{e} - O(\epsilon) \Big) g(O) - 6\epsilon d_{g,\ell},$$

assuming that ϵ is sufficiently small, say, less than $1/2$.

To see (ii), we have for $t \in \{0, \dots, 1-\epsilon\}$ that

$$\begin{aligned} L(\boldsymbol{x}^{t+\epsilon}) - L(\boldsymbol{x}^t) &= L\Big(\boldsymbol{x}^t + \epsilon \sum_{i=1}^{m} \mathbf{1}_{e_i^t} + \epsilon \boldsymbol{v}\Big) - L(\boldsymbol{x}^t) \\ &= \epsilon \sum_{i=1}^{m} \ell(e_i^t) + \epsilon(1-\epsilon)\ell(O_{\mathrm{S}}) - \epsilon^2 d_{g,\ell} && \text{(by (ii) of Lemma 4.4)} \\ &\geq \epsilon((1-\epsilon)\ell(O_{\mathrm{L}}) - \epsilon d_{g,\ell}) + \epsilon(1-\epsilon)\ell(O_{\mathrm{S}}) - \epsilon^2 d_{g,\ell} && \text{(by (ii) of Lemma 4.7)} \\ &\geq \epsilon(1-\epsilon)\ell(O) - 2\epsilon^2 d_{g,\ell}. \end{aligned}$$

By induction, we get $L(\boldsymbol{x}^1) \geq (1-\epsilon)\ell(O) - 2\epsilon d_{g,\ell}$.

To see (iii), we have for $t \in \{0, \dots, 1-\epsilon\}$ that

$$\begin{aligned} W(\boldsymbol{x}^{t+\epsilon}) - W(\boldsymbol{x}^t) &= W\Big(\boldsymbol{x}^t + \epsilon \sum_{i=1}^{m} \mathbf{1}_{e_i^t} + \epsilon \boldsymbol{v}\Big) - W(\boldsymbol{x}^t) \\ &\leq \epsilon(w(O_{\mathrm{L}}) + w(O_{\mathrm{S}})) && \text{(By (iii)'s of Lemmas 4.4 and 4.7)} \\ &= \epsilon w(O). \end{aligned}$$

By induction, we get $W(\boldsymbol{x}^1) \leq w(O)$.

Finally, we analyze the time complexity. For estimating the $\boldsymbol{\theta}_i$'s, we need $O(\frac{nm}{\epsilon} \cdot \frac{m}{\epsilon^2} \log \frac{nm}{\epsilon\delta}) = O(\frac{nm^2}{\epsilon^3} \log \frac{nm}{\epsilon\delta})$ time. The time complexity of SmallElements is at most $O(\frac{1}{\epsilon} \cdot \big(n^4 + \frac{n^2}{\epsilon^2} \log \frac{1}{\epsilon\delta}) = O(\frac{n^4}{\epsilon} + \frac{n^2}{\epsilon^3} \log \frac{1}{\epsilon\delta})$. Hence, the running time is as desired. □

4.4 Rounding

In this section, we explain how to round the vectors obtained by GuessingContinuousGreedy (Algorithm 2).

Let $(\boldsymbol{y}_1, \dots, \boldsymbol{y}_m, \boldsymbol{z})$ be the vectors obtained by GuessingContinuousGreedy, and let $\boldsymbol{v}^t$ be the vector supported on E_{S} obtained in the iteration at time t in GuessingContinuousGreedy. Note that $\boldsymbol{z} = \sum_{t \in \{0,\epsilon,\dots,1-\epsilon\}} \boldsymbol{v}^t$. Our algorithm is summarized in Algorithm 3.

We use the following lemma to analyze the objective value of the output set.

Lemma 4.9 (Lemma 3.7 of [3]). *Let $E = E_1 \cup \cdots \cup E_k$, let $f : 2^E \to \mathbb{R}_+$ be a monotone submodular function, and for all $i \neq j$, we have $E_i \cap E_j = \emptyset$. Let $\boldsymbol{x} \in \mathbb{R}_+^E$ such that for each E_i we have $\boldsymbol{x}(E_i) \leq 1$. If T is a random set where we sample independently from each E_i at most one random element, i.e., element e with probability $\boldsymbol{x}(e)$, then*

$$\mathbf{E}[f(T)] \geq F(\boldsymbol{x}).$$

Lemma 4.10. *We have $\mathbf{E}[g(S_{\mathrm{L}} \cup S_{\mathrm{S}})] \geq (1-\epsilon)\widehat{G}(\boldsymbol{x}) - \epsilon^3 v_g$ and $\mathbf{E}[\ell(S_{\mathrm{L}} \cup S_{\mathrm{S}})] \geq (1-\epsilon)L(\boldsymbol{x}) - \epsilon^3 v_\ell$.*

Algorithm 3 ROUNDING$_\epsilon(w, E_{\mathrm{L}}, E_{\mathrm{S}}, m, \{\boldsymbol{y}_i\}, \boldsymbol{z})$

Input: A weight function $w : E \to [0,1]$, $E_{\mathrm{L}}, E_{\mathrm{S}} \subseteq E$, a set of vectors $\{\boldsymbol{y}_i\}_{i\in[m]}$, and a vector $\boldsymbol{z}$.
Output: A set $S \subseteq E$.

1: $S_{\mathrm{L}} \leftarrow \emptyset$, $S_{\mathrm{S}} \leftarrow \emptyset$.
2: Define $\boldsymbol{z}' \in [0,1]^{E_{\mathrm{S}}}$ as $\boldsymbol{z}'(e) = (1-\epsilon)\boldsymbol{z}(e)$ if $w(e) < \epsilon^3 \max_t W(\boldsymbol{v}^t)$ and $\boldsymbol{z}'(e) = 0$ otherwise.
3: For each $e \in E_{\mathrm{S}}$, add it to S_{S} independently with probability $\boldsymbol{z}'(e)$.
4: **for** $i \in [m]$ **do**
5: Add exactly one element in E_i to S_{L} (as an element of E), where an element $e \in E_i$ is chosen with probability $\boldsymbol{y}_i(e)$.
6: **if** $w(S_{\mathrm{S}} \cup S_{\mathrm{L}}) \le 1$ **then**
7: **return** $S_{\mathrm{S}} \cup S_{\mathrm{L}}$.
8: **else**
9: **return** $\emptyset$.

Proof. Let $\boldsymbol{x}' = \sum_{i\in[m]} \boldsymbol{y}_i + \boldsymbol{z}'$. First, let us relate the value of the vector $\boldsymbol{x}$ to that of $\boldsymbol{x}'$.

$$\begin{aligned}\widehat{G}(\boldsymbol{x}') = \widehat{G}\Big(\sum_{i\in[m]} \boldsymbol{y}_i + \boldsymbol{z}'\Big) &\ge \widehat{G}\Big(\sum_{i\in[m]} \boldsymbol{y}_i + (1-\epsilon)\boldsymbol{z}\Big) - \sum_{e\in E_{\mathrm{S}}: w(e)\ge \epsilon^3 \max_t W(\boldsymbol{v}^t)} \boldsymbol{z}(e) g(e) \\ &\ge \widehat{G}\Big((1-\epsilon)(\sum_{i\in[m]} \boldsymbol{y}_i + \boldsymbol{z})\Big) - \max_{e\in E_{\mathrm{S}}} g(e) \sum_{e\in E_{\mathrm{S}}: w(e)\ge \epsilon^3 \max_t W(\boldsymbol{v}^t)} \boldsymbol{z}(e) \\ &\ge (1-\epsilon)\widehat{G}(\boldsymbol{x}) - \epsilon^6 v_g \cdot \frac{1}{\epsilon^3} \ge (1-\epsilon)\widehat{G}(\boldsymbol{x}) - \epsilon^3 v_g.\end{aligned}$$

Next, we note that we get S_{L} by selecting exactly one random element from each E_i, which is a copy of E_{L}, and we get S_{S} by sampling independently from $\boldsymbol{v}$. Hence, by applying Lemma 4.9 with sets $E_1, \dots, E_m$ and sets $\{\{e\} \mid e \in S_{\mathrm{S}}\}$, we get

$$\mathbf{E}[g(S_{\mathrm{L}} \cup S_{\mathrm{S}})] \ge (1-\epsilon)\widehat{G}(\boldsymbol{x}) - \epsilon^3 v_g.$$

By a similar argument, we get $L(\boldsymbol{x}') \ge (1-\epsilon)L(\boldsymbol{x}) - \epsilon^3 v_\ell$, and we have $\mathbf{E}[\ell(S_{\mathrm{L}} \cup S_{\mathrm{S}})] \ge (1-\epsilon)L(\boldsymbol{x}) - \epsilon^3 v_\ell$. □

Next, we show that the probability that the weight of the output set exceeds $w(O)$ decays exponentially.

Lemma 4.11. *For any $\gamma \ge 1$, we have $w(S_{\mathrm{L}} \cup S_{\mathrm{S}}) \le \gamma w(O)$ with probability $1 - \exp\big(-\Omega(\gamma/\epsilon^2)\big)$.*

Proof. Recall that, for each $i \in [m]$, the vector $\boldsymbol{y}_i$ is the sum of $1/\epsilon$ elements $e_i^0, \dots, e_i^{1-\epsilon}$, and we pick one of them in Algorithm 3. By the condition $w(e_i^t) \le w(o_i)$ for every $i \in [m]$ and $t \in \{0, \epsilon, \dots, 1-\epsilon\}$, the weight of the large elements after the rounding will be less than that of the large elements of the optimal solution. Hence, it is sufficient to prove that $w(S_{\mathrm{S}}) \le \gamma w(O_{\mathrm{S}})$ holds with probability $1 - \exp(-\Omega(\gamma/\epsilon^2))$, where S_{S} is the set obtained by rounding $\boldsymbol{z}'$.

First, note that

$$\mathbf{E}[w(S_{\mathrm{S}})] = \mathbf{E}[w(R(\boldsymbol{z}'))] \le (1-\epsilon)\,\mathbf{E}[w(R(\boldsymbol{z}))] \le (1-\epsilon)\max_t W(\boldsymbol{v}^t) \le (1-\epsilon)w(O_{\mathrm{S}}).$$

For each $e \in E$, we set up a random variable X_e to be $X_e = w(e)/(\epsilon^3 w(O_{\mathrm{S}}))$ if $e \in S_{\mathrm{S}}$ and $X_e = 0$ otherwise. Note that each X_e is bounded in $[0,1]$ because $\max_t W(\boldsymbol{v}^t) \le w(O_{\mathrm{S}})$. For $X = \sum_{e \in E_{\mathrm{S}}} X_e$, we have $\mu := \mathbf{E}[X] = \mathbf{E}[w(S_{\mathrm{S}})]/(\epsilon^3 w(O_{\mathrm{S}})) \le (1-\epsilon)/\epsilon^3$.

Invoking Lemma 2.2 with $\alpha = \epsilon/2$ and $\beta = \gamma/(2\epsilon^3)$, we have

$$\begin{aligned}\Pr[w(S_{\mathrm{S}}) > \gamma w(O_{\mathrm{S}})] &= \Pr\Big[X \ge \frac{\gamma}{\epsilon^3}\Big] \le \Pr\Big[X \ge (1+\alpha)\mu + \beta\Big] \\ &\le 2\exp\Big(-\frac{\alpha\beta}{3}\Big) = \exp\Big(-\Omega\Big(\frac{\gamma}{\epsilon^2}\Big)\Big).\end{aligned}$$

□

Lemma 4.12. *Algorithm 3 outputs a (random) set S with $w(S) \le 1$ satisfying*

$$\mathbf{E}[g(S) + \ell(S)] \ge (1-\epsilon)(\widehat{G}(\boldsymbol{x}) + L(\boldsymbol{x})) - O(\epsilon) \cdot (g(O) + \ell(O) + v_g + v_\ell).$$

Proof. It is clear that we always have $w(S) \le 1$.

Now, we analyze the objective value attained by S. For any $\gamma \ge 1$, the probability that $w(S_{\mathrm{L}} \cup S_{\mathrm{S}}) > \gamma w(O)$ is at most $\exp(-C\gamma/\epsilon^2)$ for some $C > 0$ by Lemma 4.11. Note that, if $T \subseteq E$ satisfies $w(T) \le \gamma w(O)$, then $g(T) + \ell(T) \le \gamma(g(O) + \ell(O))$ from the submodularity of $g + \ell$. By Lemma 4.10, we have

$$\begin{aligned}&\mathbf{E}[g(S) + \ell(S)] \\ &\ge \mathbf{E}[g(S_{\mathrm{L}} \cup S_{\mathrm{S}}) + \ell(S_{\mathrm{L}} \cup S_{\mathrm{S}})] - \int_1^\infty \gamma(g(O) + \ell(O)) \exp(-C\gamma/\epsilon^2) \mathrm{d}\gamma \\ &\ge (1-\epsilon)\widehat{G}(\boldsymbol{x}) - \epsilon^3 v_g + (1-\epsilon)L(\boldsymbol{x}) - \epsilon^3 v_\ell - \frac{\epsilon^4 + C\epsilon^2}{C^2} \exp(-C/\epsilon^2)(g(O) + \ell(O)) \\ &= (1-\epsilon)(\widehat{G}(\boldsymbol{x}) + L(\boldsymbol{x})) - O(\epsilon) \cdot (g(O) + \ell(O) + v_g + v_\ell).\end{aligned}$$

□

4.5 Putting things together

Now, we present our entire algorithm. The idea is to simply guess v_g, v_ℓ, m, $\{\gamma_i^t\}$, $\{\lambda_i\}$, $\{\gamma_{\mathrm{S}}^t\}$, and λ_{S}, run Algorithm 2 with the guessed values, and then round the obtained vectors using Algorithm 3.

Naively, we have $O(|V_{\epsilon,n}(g,\ell)|^{O(1/\epsilon)}) = O((\log(n/\epsilon)/\epsilon)^{O(1/\epsilon)})$ choices for the sequence $\{\gamma_{\mathrm{S}}^t\}$. We can decrease the number of choices since g has a bounded curvature. If we have a guess γ_{S}^0 such that $\gamma_{\mathrm{S}}^0 \ge g(O_{\mathrm{S}}) \ge (1-\epsilon)\gamma_{\mathrm{S}}^0$, then we must have $\gamma_{\mathrm{S}}^0 \ge g_S(O_{\mathrm{S}}) \ge (1-\epsilon)(1-c_g)\gamma_{\mathrm{S}}^0$ for any set $S \subseteq E$. Hence, it suffices to consider sequences whose maximum and minimum values are within a factor of $(1-\epsilon)(1-c_g)$. Let $V_{\epsilon,n,\gamma_{\mathrm{S}}^0}(g,\ell) := \{v \in V_{\epsilon,n}(g,\ell) \mid v \ge (1-\epsilon)(1-c_g)\gamma_{\mathrm{S}}^0\}$. Then, the number of such sequences is at most $|V_{\epsilon,n}(g,\ell)| \cdot |V_{\epsilon,n,\gamma_{\mathrm{S}}^0}(g,\ell)|^{O(1/\epsilon)}$, which is much smaller than $O((\log(n/\epsilon)/\epsilon)^{O(1/\epsilon)})$.

A detailed description of our algorithm is given in Algorithm 4.

Proof of Theorem 3.1. Consider the case that v_g and v_ℓ satisfy (2), $m = |O_{\mathrm{L}}|$, and $\{\gamma_i^t\}, \{\lambda_i\}, \{\gamma_{\mathrm{S}}^t\}$, and λ_{S} are good guesses. Let S be the (random) set obtained with these guesses. By Lemma 4.12, we have

$$\mathbf{E}[g(S) + \ell(S)] \ge (1-\epsilon)(\widehat{G}(\boldsymbol{x}) + L(\boldsymbol{x})) - O(\epsilon)(g(O) + \ell(O) + v_g + v_\ell). \tag{3}$$

Algorithm 4 KNAPSACK

Input: A monotone submodular function $g : 2^E \to \mathbb{R}_+$, a linear function $\ell : 2^E \to \mathbb{R}$, a weight function $w : E \to \mathbb{R}_+$, and $\epsilon \in (0, 1)$.
Output: A set $S \subseteq E$ satisfying $w(S) \leq 1$.

for each choice of $v_g, v_\ell \in V_{\epsilon,n}(g, \ell)$ **do**
  $E_{\mathrm{L}} \leftarrow$ the set of large elements with respect to v_g and v_ℓ.
  $E_{\mathrm{S}} \leftarrow$ the set of small elements with respect to v_g and v_ℓ.
  $\mathcal{S} \leftarrow \emptyset$.
  $M := \lfloor \frac{1}{(1-c_g)\epsilon^6} \rfloor$
  for each choice of m from $\{0, 1, \ldots, M\}$ **do**
    for each choice of $\{\gamma_i^t\}, \{\lambda_i\}$ from $V_{\epsilon,m}(g, \ell)/m$ **do**
      for each choice of $\gamma_{\mathrm{S}}^0, \lambda_{\mathrm{S}}$ from $V_{\epsilon,n}(g, \ell)$ **do**
        for each choice of $\{\gamma_{\mathrm{S}}^\epsilon, \ldots, \gamma_{\mathrm{S}}^{1-\epsilon}\}$ from $V_{\epsilon,n,\gamma_{\mathrm{S}}^0}(g, \ell)$ **do**
          $(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m, \boldsymbol{z}) \leftarrow \textsc{GuessingContinuousGreedy}_{\epsilon,\epsilon}(g, \ell, w, E_{\mathrm{L}}, E_{\mathrm{S}}, m, \{\gamma_i^t\}, \{\lambda_i\}, \{\gamma_{\mathrm{S}}^t\}, \lambda_{\mathrm{S}})$.
          $S \leftarrow \textsc{Rounding}_\epsilon(w, E_{\mathrm{L}}, E_{\mathrm{S}}, m, \{\boldsymbol{y}_i\}, \boldsymbol{z})$.
          $\mathcal{S} \leftarrow \mathcal{S} \cup \{S\}$.
return $\arg\max_{S \in \mathcal{S}} g(S) + \ell(S)$.

Conditioned on the event that GuessingContinuousGreedy succeeds, by (i) and (ii) of Lemma 4.8, we get

$$\begin{aligned}(3) &\geq (1-\epsilon)(1-1/e-O(\epsilon))g(O) + (1-\epsilon)(1-O(\epsilon))\ell(O) - O(\epsilon)(g(O)+\ell(O)+v_g+v_\ell+8d_{g,\ell}) \\ &\geq (1-1/e)g(O) + \ell(O) - O(\epsilon)(g(O)+\ell(O)). \qquad (4)\end{aligned}$$

Since GuessingContinuousGreedy succeeds with probability at least $1-\epsilon$, we get

$$\mathbf{E}[g(S)+\ell(S)] \geq (1-\epsilon)\cdot(4) \geq (1-1/e)g(O) + \ell(O) - O(\epsilon)(g(O)+\ell(O)).$$

Since Algorithm 4 outputs the set with the maximum objective, we have the desired property on the objective value.

It is clear that the output of Algorithm 4 has weight at most 1 because Rounding always outputs a set of weight at most 1.

For arbitrary $\gamma \in V_{\epsilon,n}(g, \ell)$, the time complexity of Algorithm 4 is

$$\begin{aligned}&O\Big(\frac{nM^2}{\epsilon^3}\log\frac{nM}{\epsilon} + \frac{n^4}{\epsilon} + \frac{n^2}{\epsilon^3}\log\frac{1}{\epsilon}\Big)\cdot|V_{\epsilon,n}(g,\ell)|^{O(1)}\cdot|V_{\epsilon,n,\gamma}(g,\ell)|^{O(1/\epsilon)}\cdot|V_{\epsilon,m}(g,\ell)|^{O(M/\epsilon)} \\ &= O\Big(\frac{nM^2}{\epsilon^3}\log\frac{nM}{\epsilon} + \frac{n^4}{\epsilon} + \frac{n^2}{\epsilon^3}\log\frac{1}{\epsilon}\Big)\cdot\Big(\frac{\log(n/\epsilon)}{\epsilon}\Big)^{O(1)}\cdot\Big(\frac{1}{\epsilon}\log\frac{1}{1-c_g}\Big)^{O(1/\epsilon)}\cdot\Big(\frac{\log(M/\epsilon)}{\epsilon}\Big)^{O(M/\epsilon)} \\ &= O\Big(\frac{n}{(1-c_g)^2\epsilon^{15}}\log\frac{n}{(1-c_g)\epsilon} + \frac{n^4}{\epsilon} + \frac{n^2}{\epsilon^3}\log\frac{1}{\epsilon}\Big)\cdot\Big(\frac{\log n}{\epsilon}\Big)^{O(1)}\cdot \\ &\quad \Big(\frac{1}{\epsilon}\log\frac{1}{1-c_g}\Big)^{O(1/\epsilon)}\cdot\Big(\frac{1}{\epsilon}\log\frac{1}{1-c_g}\Big)^{O(1/((1-c_g)\epsilon^7))} \\ &= O\Big(\frac{n^4\mathrm{polylog}(n)}{(1-c_g)^2}\Big)\cdot\Big(\frac{1}{\epsilon}\log\frac{1}{1-c_g}\Big)^{\mathrm{poly}(1/\epsilon)/(1-c_g)}.\end{aligned}$$

Hence, we have the desired time complexity.

By replacing ϵ with ϵ/C for a large constant C (to change $O(\epsilon)$ to ϵ), we have the desired result. □

5 The Budget Allocation Problem

In this section, we bound the curvature of the submodular function that represents the budget allocation problem, and we confirm that our algorithm can be applied to the budget allocation problem in order to obtain an approximation factor better than $1 - 1/e$.

We formally define the budget allocation problem. The input consists of a bipartite graph with the bipartition $A \cup B$, a weight function $w : A \to [0, 1]$, a capacity function $c : A \to \mathbb{N}$, and a probability function $p : A \to [0, 1]$. Intuitively speaking, the sets A and B correspond to media channels and customers, respectively. Each edge (a, b) in the bipartite graph represents the potential influence of media channel a on customer b. Consider a budget allocation $\boldsymbol{b} \in \mathbb{Z}_+^A$ to A with $\boldsymbol{b}(a) \leq c(a)$ and $\sum_{a \in A} \boldsymbol{b}(a)w(s) \leq 1$. If a node a is allocated a budget of $\boldsymbol{b}(a)$, it makes $\boldsymbol{b}(a)$ independent trials to activate each adjacent node b. The probability that b is activated by a in each trial is $p(a)$. Thus, the probability that b becomes active is

$$1 - \prod_{a \in \Gamma(b)} p(a)^{\boldsymbol{b}(a)},$$

where $\Gamma(b)$ denotes the set of nodes in A adjacent to b. Hence, the expected number of activated target nodes is

$$\sum_{b \in B} \Big(1 - \prod_{a \in \Gamma(b)} p(a)^{\boldsymbol{b}(a)}\Big).$$

The objective of this problem is to find the budget allocation that maximizes the expected number of activated target nodes.

We can recast the problem using a submodular function. For each $a \in A$, let $E_a = \{(a, i) \mid i \in c(a)\}$, and let $E = \bigcup_{a \in A} E_a$. Then, we define $f : 2^E \to \mathbb{R}_+$ as

$$f(S) = \sum_{b \in B} \Big(1 - \prod_{a \in \Gamma(b)} p(a)^{|S \cap E_a|}\Big).$$

Further, we define $w' : 2^E \to [0, 1]$ to be $w'((a, i)) = w(a)$. Then, the budget allocation problem is equivalent to maximizing $f(S)$ subject to $w'(S) \leq 1$.

We now observe several properties of f.

Lemma 5.1. *Let $S \subsetneq E$ and $(a, i) \in E \setminus S$. Then,*

$$f_S((a, i)) = \sum_{b \in B : a \in \Gamma(b)} (1 - p(a)) \prod_{a' \in \Gamma(b)} p(a')^{|S \cap E_{a'}|}.$$

Proof. For each $b \in B$, we define a function $g^b : 2^E \to \mathbb{R}_+$ as $g^b(T) = 1 - \prod_{a \in \Gamma(b)} p(a)^{|T \cap E_a|}$. Note that $f_S((a, i)) = \sum_{b \in B} g^b_S((a, i))$.

If $a \notin \Gamma(b)$, then we clearly have $g^b_S((a, i)) = 0$. If $a \in \Gamma(b)$, then we have

$$g^b_S((a, i)) = (1 - p(a)) \prod_{a' \in \Gamma(b)} p(a')^{|S \cap E_{a'}|}.$$

Summing $g_S^b((a,i))$ over all $b \in B$, we obtain the claim. □

Corollary 5.2. *The function f is submodular.*

Proof. From Lemma 5.1, it is easy to see that $f_S((a,i)) \geq f_T((a,i))$ holds for $S \subseteq T \subsetneq E$ and $(a,i) \in E \setminus T$. □

Corollary 5.3. *The curvature c_f of f satisfies*

$$c_f \leq 1 - \min_{a \in A} \min_{b \in B : a \in \Gamma(b)} p(a)^{c(a)-1} \prod_{a' \in \Gamma(b) \setminus \{a\}} p(a')^{c(a')}.$$

Proof. From Lemma 5.1, we have

$$\begin{aligned} f_{E \setminus (a,i)}((a,i)) &= \sum_{b \in B : a \in \Gamma(b)} (1 - p(a)) p(a)^{c(a)-1} \prod_{a' \in \Gamma(b) \setminus \{a\}} p(a')^{c(a')}, \\ f((a,i)) &= \sum_{b \in B : a \in \Gamma(b)} (1 - p(a)). \end{aligned}$$

Hence,

$$\begin{aligned} c_f &= 1 - \min_{(a,i) \in E} \frac{f_{E \setminus (a,i)}((a,i))}{f((a,i))} \\ &= 1 - \min_{a \in A} \frac{\sum_{b \in B : a \in \Gamma(b)} (1 - p(a)) p(a)^{c(a)-1} \prod_{a' \in \Gamma(b) \setminus \{a\}} p(a')^{c(a')}}{\sum_{b \in B : a \in \Gamma(b)} (1 - p(a))} \\ &\leq 1 - \min_{a \in A} \min_{b \in B : a \in \Gamma(b)} p(a)^{c(a)-1} \prod_{a' \in \Gamma(b) \setminus \{a\}} p(a')^{c(a')}. \end{aligned}$$

□

From our main result (Theorem 1.1) and Corollaries 5.2 and 5.3, when the capacity of each node $a \in A$ is bounded by a constant and the number of vertices adjacent to each node $b \in B$ is bounded by a constant, we obtain a polynomial-time algorithm whose approximation ratio is strictly better than $1 - 1/e$.

Acknowledgments

We thank Takanori Maehara for providing us with the problem and insightful comments.

References

[1] N. Alon, I. Gamzu, and M. Tennenholtz. Optimizing budget allocation among channels and influencers. In *Proceedings of the 21st International Conference on World Wide Web (WWW)*, pages 381–388, 2012.

[2] A. Badanidiyuru and J. Vondrák. Fast algorithms for maximizing submodular functions. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1497–1514, 2013.

[3] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

[4] M. Conforti and G. Cornuéjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete Applied Mathematics*, 7(3):251–274, 1984.

[5] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.

[6] M. Feldman. *Maximization Problems with Submodular Objective Functions*. PhD thesis, Technion, 2013.

[7] R. K. Iyer and J. A. Bilmes. Submodular optimization with submodular cover and submodular knapsack constraints. In *Proceedings of the 27th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2436–2444, 2013.

[8] A. Krause, A. P. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008.

[9] A. Kulik, H. Shachnai, and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 545–554, 2013.

[10] J. Lee. *Maximum Entropy Sampling*. John Wiley & Sons, Ltd, 2006.

[11] H. Lin and J. Bilmes. Multi-document summarization via budgeted maximization of submodular functions. pages 912–920. Proceedings of the 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 2010.

[12] H. Lin and J. Bilmes. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT)*, pages 510–520, 2011.

[13] D. Sharma, A. Kapoor, and A. Deshpande. On greedy maximization of entropy. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1330–1338, 2015.

[14] T. Soma, N. Kakimura, K. Inaba, and K. Kawarabayashi. Optimal budget allocation: Theoretical guarantee and efficient algorithm. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 351–359, 2014.

[15] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.

[16] M. Sviridenko, J. Vondrák, and J. Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1134–1148, 2015.

[17] J. Vondrák. Submodularity and curvature: the optimal algorithm. *RIMS Kokyuroku Bessatsu*, B23:253–266, 2010.